arXiv:1609.08251v1 [math.NA] 27 Sep 2016

Robust and efficient multi-way spectral clustering

Anil Damle*, Victor Minden†, and Lexing Ying‡

September 28, 2016

Abstract

We present a new algorithm for spectral clustering based on a column-pivoted QR factorization that may be directly used for cluster assignment or to provide an initial guess for `k-means`. Our algorithm is simple to implement, direct, and requires no initial guess. Furthermore, it scales linearly in the number of nodes of the graph and a randomized variant provides significant computational gains. Provided the subspace spanned by the eigenvectors used for clustering contains a basis that resembles the set of indicator vectors on the clusters, we prove that both our deterministic and randomized algorithms recover a basis close to the indicators in Frobenius norm. Finally, we experimentally demonstrate that the performance of our algorithm tracks recent information theoretic bounds for exact recovery in the stochastic block model.

Spectral clustering has found extensive use as a mechanism for detecting well-connected subgraphs of a network. Typically, this procedure involves computing an appropriate number of eigenvectors of the (normalized) Laplacian and subsequently applying a clustering algorithm to the embedding of the nodes defined by the eigenvectors. Currently, one of the most popular algorithms is `k-means++` [5], the standard iterative `k-means` algorithm [20] applied to an initial clustering chosen via a specified random sampling procedure. Due to the non-convex nature of the `k-means` objective, however, this initialization does not preclude convergence to local minima, which can be poor clusterings.

We provide an alternative, direct (non-iterative) procedure for clustering the nodes in their eigenvector embedding. It is important to note that our procedure is not a substitute for `k-means++` when tackling general (*i.e.*, non-spectral) clustering problems. For spectral embeddings of graphs with community structure, however, we take advantage of additional geometric structure of the embedding to build a more robust clustering procedure. Furthermore, our algorithm is built out of a simple column-pivoted QR factorization, making it easy to implement and use. Finally, a simple randomized acceleration of our algorithm substantially reduces the cost of cluster assignment, making it feasible for large problems of practical interest.

1 Background and setup

Given a simple undirected graph $\mathcal{G}$ with adjacency matrix $A \in \{0,1\}^{n\times n}$, we consider the multi-way clustering problem of partitioning the vertices of $\mathcal{G}$ into k disjoint clusters. A

*Department of Mathematics, University of California, Berkeley (damle@berkeley.edu)

†Institute for Computational & Mathematical Engineering, Stanford University (vminden@stanford.edu)

‡Department of Mathematics and Institute for Computational & Mathematical Engineering, Stanford University (lexing@stanford.edu)

common (albeit unrealistic) generative model for graphs exhibiting this sort of cluster structure is is the *k-way stochastic block model*.

Definition 1 (Stochastic block model [18])**.** *Partition $[n]$ into k mutually disjoint and non-empty clusters $\mathcal{C}_1, \ldots, \mathcal{C}_k$. Given probabilities p and q such that $p > q$, let $M \in [0,1]^{n \times n}$ have entries $M_{ii} = 0$ and, for $i \neq j$,*

$$M_{ij} = \begin{cases} p, & \{i,j\} \subset \mathcal{C}_k \text{ for some } k, \\ q, & \text{else.} \end{cases}$$

A symmetric adjacency matrix $A \in \{0,1\}^{n \times n}$ with $A_{ij} \sim \text{Bernoulli}(M_{ij})$ for $i < j$ and $A_{ii} = 0$ for all i is said to be distributed according to the k-way stochastic block model (SBM) with clusters $\{\mathcal{C}_i\}_{i=1}^k$, within-cluster probability p, and between-cluster probability q.

For an SBM with equisized clusters, the maximum-likelihood estimate for the clusters can be phrased in terms of maximizing the number of within-cluster edges; that is, given A, find a matrix X whose columns are indicator vectors for cluster membership such that X attains the optimal value of the combinatorial optimization problem

$$\begin{aligned} \underset{X}{\text{maximize}} \quad & \operatorname{Tr}\left(X^T A X\right) \\ \text{subject to} \quad & X \in \{0,1\}^{n \times k}, \quad X^T X = \frac{n}{k} I_k. \end{aligned} \tag{1}$$

If A is not assumed to be a random sample from the SBM, then the above problem does not have the interpretation of maximum-likelihood estimation, though it remains a common starting point for clustering.

Given that this combinatorial optimization problem is NP-hard, it is typical to relax (1) to a computationally-tractable convex formulation. One major avenue of work in recent years has been in the area of semidefinite programming (SDP) relaxation for clustering [17, 2, 1]. Broadly speaking, such relaxations recast (1) in terms of XX^T and then relax XX^T to a semidefinite matrix Z. These SDP relaxations often enjoy strong consistency results on recovery down to the information theoretic limit in the case of the SBM, in which setting the optimal solution Z_* can be used to recover the true clusters exactly with high probability.

A more common relaxation of (1) is to remove the restriction that $X \in \{0,1\}^{n \times k}$ and instead optimize over real-valued matrices,

$$\begin{aligned} \underset{X}{\text{maximize}} \quad & \operatorname{Tr}\left(X^T A X\right) \\ \text{subject to} \quad & X \in \mathbf{R}^{n \times k}, \quad X^T X = \frac{n}{k} I_k. \end{aligned} \tag{2}$$

While this optimization problem is still non-convex, it follows from the Courant-Fischer-Weyl min-max principle that an optimal point X_* is given as $X_* = V_k Q$, where $V_k \in \mathbf{R}^{n \times k}$ contains the eigenvectors of A corresponding to the k largest eigenvalues and $Q \in \mathbf{O}_k$ is an arbitrary orthogonal transformation. Because the solution X_* is no longer discrete, the canonical spectral clustering approach uses the rows of X_* as coordinates in a standard point-cloud clustering procedure such as `k-means`.

We propose an algorithm based on a column-pivoted QR factorization of the matrix V_k^T that can be used either as a stand-alone clustering algorithm or to initialize iterative algorithms such as `k-means`. Our approach stems from the computational quantum chemistry

literature, where the same basic procedure has been developed as a method for constructing spatially-localized bases of an eigenspace associated with the nonlinear eigenproblem of Kohn-Sham density functional theory [10, 9]. Our numerical experiments show that our approach closely tracks the information-theoretic limit in terms of exact clustering of the stochastic block model, whereas `k-means++` on its own does not.

1.1 Related work

A similar algorithm using a column-pivoted QR for cluster assignment of general point-cloud data has appeared in the literature previously in the context of spectral relaxation of the `k-means` objective [31]. Curiously, though we find the basic idea to be powerful, this algorithmic approach seems to have been ignored and we can find no further reference to it. We expand upon this work, taking advantage of algorithmic improvements appearing in the computational chemistry literature for greater efficiency. Further, by addressing the sparse adjacency matrix case instead of the Gram matrix of point-cloud data, we are able to strongly motivate our approach based on proven spectral properties of model graphs.

Recently, there has been significant work on understanding when it is possible to exactly recover communities in the sparse SBM, wherein the probability of connections between nodes is $\Theta(\log n/n)$. Specifically, the information theoretic bound for when exact recovery is possible with two blocks was developed by Abbe et al. [1] and an SDP-based algorithm achieving the bound was proposed. Additional recent work [2, 16, 17] has extended this theory to the multi-block case (for a constant or slowly growing number of clusters) and generalized the SDP approach.

Spectral methods as we discuss here stem from work on the Fiedler vector [13, 12] for the two-block case and spectral embeddings coupled with `k-means` clustering [25] in the multi-block case. For a more comprehensive overview of initialization procedures for `k-means` see, *e.g.* , [7]. Recent analysis of these methods applied to the SBM [27] focuses on misclassification of nodes. Another line of work considers matrices besides the adjacency matrix or normalized Laplacian to achieve theoretical detection thresholds [19, 22]. Other recent work [23] demonstrates where spectral methods break down and argues for the use of SDP-based methods.

2 Algorithm

2.1 Preliminaries

Suppose $\mathcal{G}$ is a simple undirected graph with symmetric adjacency matrix $A \in \{0,1\}^{n\times n}$. Let the eigendecomposition of A be given by

$$A = \left[\begin{array}{c|c} V_k & V_{n-k} \end{array}\right] \left[\begin{array}{c|c} \Lambda_k & 0 \\ \hline 0 & \Lambda_{n-k} \end{array}\right] \left[\begin{array}{c|c} V_k & V_{n-k} \end{array}\right]^T,$$

where $V_k \in \mathbf{R}^{n\times k}$ and $V_{n-k} \in \mathbf{R}^{n\times(n-k)}$ contain pairwise orthonormal columns of eigenvectors, and $\Lambda_k = \mathrm{diag}(\lambda_1, \lambda_2, \ldots, \lambda_k)$ contains the eigenvalues sorted as $\lambda_1 \geq \lambda_2, \geq \cdots \geq \lambda_k$. Of interest in spectral clustering is the structure of the eigenspace spanned by the columns of V_k or a related matrix, as we make explicit below.

Definition 2. *Suppose $\mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_k$ is a clustering of the vertex set $[n]$. We say the matrix $W' \in \{0,1\}^{n\times k}$ with entries given by*

$$W'_{ij} = \begin{cases} 1, & \text{if vertex } i \text{ belongs to cluster } j, \\ 0, & \text{otherwise,} \end{cases}$$

is an indicator matrix *for the underlying clustering. If W' is an indicator matrix and $W \in \mathbf{R}_+^{n\times k}$ is given by scaling the columns of W' to have unit ℓ_2 norm then we call W a* normalized indicator matrix *for the underlying clustering.*

Suppose A is sampled from the k-way SBM on n nodes with within- and between-cluster connection probabilities p and q. Letting W be a normalized indicator matrix for the underlying clustering and provided that p and q do not decay too quickly asymptotically, standard matrix concentration inequalities [29] coupled with the Davis-Kahan theorem [11] give convergence of the spectral projector $V_k V_k^T$ in n, *i.e.*, for fixed k there exists $\delta_n = o(1)$ such that

$$\left\|WW^T - V_k V_k^T\right\|_2 \leq \delta_n$$

with high probability. The exact nature of these concentration results depends strongly on the asymptotic behavior of p and q, however, the general idea is that with high probability any given A will be close to the average adjacency matrix. More generally, if A is a symmetric adjacency matrix not necessarily coming from the SBM that is sufficiently close to M, the same basic argument applies [30]. Therefore, let $\epsilon > 0$ be such that

$$\left\|WW^T - V_k V_k^T\right\|_2 \leq \epsilon.$$

From this bound on the distance between spectral projectors we can conclude the existence of $Q \in \mathbf{O}_k$ such that

$$\|W - V_k Q\|_F \leq \sqrt{k}\epsilon + \mathcal{O}(\epsilon^2). \tag{3}$$

Concisely, we expect the eigenvector matrix V_k to look like the normalized indicator matrix W, under a suitable orthogonal transformation. Alternatively, $V_k Q$ provides an approximately sparse basis for the range of V_k. Thus, while any orthogonal transformation of the (appropriately scaled) eigenvectors is a solution to (2), there exists a specific solution that directly reflects the underlying cluster structure.

If W and V_k in (3) are both known, finding an orthogonal transformation to make V_k resemble W is not a difficult problem. To find this transformation without knowing the clustering a priori, however, we use a procedure based on the column-pivoted QR factorization. We define this factorization notationally for matrices that are wider than they are tall, which is the setting we will require later.

Definition 3. *Let $B \in \mathbf{R}^{k\times n}$ with $k \leq n$ have rank k. The* column-pivoted QR factorization *(CPQR) of B is the factorization*

$$B\Pi = Q\begin{bmatrix} R_1 & R_2 \end{bmatrix}$$

as computed via the Golub-Businger algorithm [6] (see also Chapter 5 of Golub [14]) where $\Pi \in \mathbf{R}^{n\times n}$ is a permutation matrix, $Q \in \mathbf{O}_k$ is an orthogonal matrix, $R_1 \in \mathbf{R}^{k\times k}$ is upper triangular, and $R_2 \in \mathbf{R}^{k\times(n-k)}$.

Ideally, the permutation Π is chosen to make R_1 as well-conditioned as possible given A. The Golub-Businger algorithm chooses this permutation using a greedy heuristic wherein at each step of the algorithm the column with largest norm is picked as the pivot and the remaining columns are orthogonalized against it.

2.2 Algorithm statement

Given V_k we propose the following algorithm for finding k clusters in the graph:

1. Compute the CPQR factorization
$$V_k^T \Pi = QR.$$
2. For each $j \in [n]$ assign node j to cluster
$$c_j = \arg\max_i \left(\left| R\Pi^T \right|_{i,j} \right).$$

Note that $|\cdot|$ is used to represent the elementwise absolute value operation when applied to a vector and here we assume that every node belongs to exactly one cluster. For real-world applications, if all entries in a column are small we may optionally not assign the corresponding node to any cluster, but we do not explore that further at this time.

The above algorithm is strikingly simple, there is no dependence on any sort of initial guess, it is naturally invariant to the data ordering, and it is constructed out of a single common and efficient matrix factorization. From an implementation standpoint this means we may simply use, *e.g.* , the `qr` function in MATLAB® or the LAPACK [4] routine `DGEQP3` [26]. In fact, the cost of this algorithm is linear in the number of nodes of the graph and quadratic in the number of clusters, *i.e.* $\mathcal{O}(nk^2)$. The lack of an initial guess is a particularly important distinction from the generally used `k-means` algorithm.

A simple randomized variant of our algorithm allows us to more efficiently compute a matrix Q such that $Q^T V_k$ may be used for cluster assignment with high probability. To do this, we define the probability distribution ρ over $[n]$ as
$$\Pr(\{j\}) = \frac{\left\| \left(V_k^T\right)_{:,j} \right\|_2}{k}.$$
Note that this corresponds to the so-called leverage scores [21] of the spectral projector $V_k V_k^T$. Now, given V_k, an oversampling factor γ, and a failure probability δ, the following algorithm computes the cluster assignments:

1. Sample at least $\gamma k \log\left(\frac{k}{\delta}\right)$ nodes from $[n]$ based on the discrete distribution ρ and denote this set $\mathcal{J}$.
2. Compute the CPQR factorization
$$\left(V_k^T\right)_{:,\mathcal{J}} \Pi = QR.$$
3. For each $j \in [n]$ assign node j to cluster
$$c_j = \arg\max_i \left(\left| Q^T V_k^T \right|_{i,j} \right).$$

The cost of the CPQR is now $\mathcal{O}(k^3 \log k)$, yielding a cluster assignment operator in time independent of n. While computing $Q^T V_k$ in step 3 formally scales as $\mathcal{O}(nk^2)$, the dominant cost in the deterministic algorithm is the CPQR. If one only wishes to cluster a subset of the nodes or check if two nodes are in the same cluster, then Q^T need only be applied to some columns of V_k^T. Furthermore, $Q^T V_k^T$ scales well when computed on parallel architectures and the assumed structure, specifically the locality, of the product could also potentially be used to reduce the cost of its computation. The other dependence on n in the form of $\mathcal{O}(nk)$ is in computing the column norms of V_k^T to build ρ. This operation also parallelizes well and asymptotically costs no more than writing down V_k itself. To truly make the cost of constructing Q independent of n, however, one may sample nodes uniformly to build $\mathcal{J}$. In this case one must sample slightly more than $k \log k$, in a manner that depends on the size of the smallest cluster relative to n, to ensure the same failure probability.

Prior to embarking on the proofs, we provide a brief high-level rationale for the use of a CPQR. Assuming every point belongs to exactly one cluster, the `k-means` algorithm applied to the spectral embedding may be expected to work in principle given a sufficiently close initial guess for the cluster centers. It turns out the the CPQR procedure achieves this goal when we are working with the spectral embedding for a graph with community structure.

Heuristically, consider picking an initial node at random. It belongs to one of the clusters and thus might be a reasonable surrogate for a cluster center that represents that cluster. Now, we need a procedure to pick a point that is guaranteed to be in a different cluster to use as another center. For general clustering problems this may be difficult. However, here we observe that the clusters have additional structure: we expect them to be orthogonal. This motivates the use of the CPQR. Once we orthogonalize all of the points against the center point of the first cluster, the norms of other points in the same cluster become small while the norms of points in different clusters remain more or less the same. Therefore, picking the point with largest remaining norm should yield a viable center for a different cluster. By repeating this procedure k times we hope to get a good set of k different cluster centers. This procedure is codified by the CPQR factorization introduced earlier.

2.3 Analysis

Theorem 7.2 of Gu & Eisenstat [15] proves that for a CPQR as defined above

$$\sigma_k(R_1) \geq \frac{\sigma_k(B)}{f(n,k)}, \tag{4}$$

where σ_j denotes the jth singular value of a matrix and $f(n,k) = 2^k\sqrt{n}$. This bound ensures that the singular values of R_1 are not too much smaller than the singular values of B. In fact, the CPQR defined here is related to a broader class of rank-revealing QR (RRQR) factorizations [8, 15] that seek a permutation Π such that (4) holds for different forms of the denominator $f(n,k)$, *e.g.* $\sqrt{k(n-k) + \min(n, n-k)}$. Because the matrices we deal with are quite far from the pathological worst-case examples we find the standard CPQR algorithm sufficient for our needs.

Since $R\Pi^T = Q^T V_k^T$ we will show our algorithm finds a matrix Q such that $V_k Q$ is close to W (up to a signed column permutation) even though it knows nothing explicit about the sparsity structure of W. In fact the column order in the permutation simply corresponds

to the relabeling of the clusters. All we need to do is assign each node to a cluster based on its largest magnitude entry.

We now prove two lemmas that ensures the CPQR pivoting strategy applied to V_k^T or a randomly selected subset of its columns implicitly identifies k linearly independent columns of WW^T.

Lemma 1. *Let $W \in \mathbf{R}^{n\times k}$ and $V \in \mathbf{R}^{n\times k}$ have orthonormal columns and satisfy*

$$\left\|WW^T - VV^T\right\|_2 \leq \epsilon$$

with $\epsilon < 2^{-k}/\sqrt{n}$. If $V^T\Pi = QR$ is a CPQR factorization and $\mathcal{C}$ denotes the original indices of the first k columns selected by the permutation Π, then

$$\text{range}\left\{\left(WW^T\right)_{:,\mathcal{C}}\right\} = \text{range}\left\{W\right\}.$$

Proof. Theorem 7.2 of Gu & Eisenstat [15] coupled with the fact that $\sigma_i(V) = 1$ for $i = 1,\ldots,k$ implies that $\sigma_{\min}(R_1) \geq 2^{-k}/\sqrt{n}$. Substituting in the CPQR factorization yields

$$\left(WW^T + E\right)\Pi = VQ\begin{bmatrix}R_1 & R_2\end{bmatrix}$$

with $\|E\|_2 \leq \epsilon$. Now, $\sigma_{\min}(R_1) \geq \epsilon$ implies that $\left(W^T\right)_{:,\mathcal{C}}$ is non-singular, as otherwise the distance from VQR_1 to the nearest singular matrix is less than ϵ, which is a contradiction. The result follows. □

Lemma 2. *Let $W \in \mathbf{R}^{n\times k}$ be a normalized indicator matrix and $V \in \mathbf{R}^{n\times k}$ have orthonormal columns and satisfy*

$$\left\|WW^T - VV^T\right\|_2 \leq \epsilon$$

with $\epsilon < 2^{-k}/\sqrt{n}$ and assume there exists $\gamma > 0$ such that

$$\frac{1}{\gamma}\left\|W_{:,i}\right\|_2^2 \leq \left\|V_{:,i}\right\|_2^2$$

for $i \in [n]$. Let $\mathcal{J}$ denote $\gamma k \log\frac{k}{\delta}$ samples with replacement from the discrete distribution ρ over $[n]$. If $\left(V^T\right)_{:,\mathcal{J}}\Pi = QR$ is a CPQR factorization and $\mathcal{C}$ denotes the original indices of the first k columns selected by the permutation Π, then with probability $1-\delta$

$$\text{range}\left\{\left(WW^T\right)_{:,\mathcal{C}}\right\} = \text{range}\left\{W\right\}.$$

Proof. The matrix WW^T has k distinct linearly independent columns, and because of the normalization the sum of square column norms for each set is 1. Therefore, the distribution ρ has mass $\frac{1}{k}$ on each set of linearly independent columns. Now, a simple coupon-collecting argument [24] over k bins where the smallest mass a bin can have is $1/\gamma k$ ensures that $\mathcal{J}$ contains k linearly independent columns. The result now follows immediately from Lemma 1. □

Note that the assumption of a multiplicative bound on the squared column norms can be relaxed to a bound on the sum of the squared column norms of columns associated with each cluster and the result still holds. Further, one may substitute the CPQR for a different

RRQR [8] or strong RRQR [15] in either lemma and inherit the appropriate upper bound on ϵ based on the corresponding $f(n,k)$.

The key aspect of Lemma 1 and Lemma 2 for our purposes is that the CPQR correctly identifies the subspace spanned by W. Using our additional structural assumptions we show that if this permutation had been applied *a priori* a traditional QR factorization achieves our goal. While the QR factorization is not computed this way and the permutation is computed instead at the same time as the Q and R factors, the two paths are equivalent from an analysis standpoint.

Theorem 1. *Let $W \in \mathbf{R}^{n\times k}$ be a normalized indicator matrix and suppose $n > 4$. Let $V \in \mathbf{R}^{n\times k}$ with orthonormal columns satisfy*

$$\left\|WW^T - VV^T\right\|_2 \leq \epsilon$$

with $\epsilon < 2^{-k}/\sqrt{n}$. If $V^T\Pi = QR$ is a CPQR factorization of V^T, then there exists a permutation matrix $\widehat{\Pi}$ and a diagonal matrix D with $D_{ii} = \pm 1$ such that

$$\|W\widehat{\Pi}D - VQ\|_F \leq \epsilon(1+\sqrt{2})k\sqrt{n} + \mathcal{O}\left(\epsilon^2\right).$$

Proof. Let $\mathcal{C}$ denote the original indices of the first k columns selected by the permutation Π. Based on the nonzero structure of W and Lemma 1, there exists a permutation matrix $\widehat{\Pi}$ such that $R_I \equiv \widehat{\Pi}^T \left(W^T\right)_{:,\mathcal{C}}$ is diagonal with positive diagonal entries.

From our assumptions on W and V it follows that

$$\left\|W\widehat{\Pi}R_I - VQR_1\right\|_2 \leq \epsilon$$

where the upper triangular matrix R_1 is the left-most $k \times k$ block of R. Introducing the diagonal matrix D with $D_{ii} = \pm 1$ such that $\widehat{R}_1 \equiv R_1 D$ has positive diagonal entries, we may write

$$\left\|W\widehat{\Pi}DR_I - VQ\widehat{R}_1\right\|_F \leq k\epsilon$$

using the standard two-norm bound on the Frobenius norm. Because $W\widehat{\Pi}D$ has orthonormal columns and R_I is diagonal, this is a QR factorization of some matrix M. Similarly, VQ and $\widehat{R}_1$, together form the QR factorization of a matrix that may be written using the above bound as $M + E$ for some E with $\|E\|_F \leq k\epsilon$.

Given that R_I and $\widehat{R}_1$ have positive diagonal entries and the technical condition $k\epsilon \leq 1/4$ (which follows from our assumptions), we may now use a result on perturbations of QR factorizations from Eq. (3.2) of Stewart [28] to conclude that

$$\|W\widehat{\Pi}D - VQ\|_F \leq \kappa_2(R_I)\frac{\left(1+\sqrt{2}\right)\|E\|_F}{\|R_I\|_2} + \mathcal{O}\left(\epsilon^2\right).$$

The conclusion of the theorem follows from the fact that $\sigma_k(R_I) \geq 1/\sqrt{n}$, as this is a lower bound on the entries of a $n \times k$ normalized indicator matrix. □

Theorem 2. *Let $W \in \mathbf{R}^{n\times k}$ be a normalized indicator matrix and suppose $n > 4$. Let $V \in \mathbf{R}^{n\times k}$ with orthonormal columns satisfy*

$$\left\|WW^T - VV^T\right\|_2 \leq \epsilon$$

with $\epsilon < 2^{-k}/\sqrt{n}$ *and assume there exists* $\gamma > 0$ *such that*

$$\frac{1}{\gamma} \|W_{:,i}\|_2^2 \leq \|V_{:,i}\|_2^2$$

for $i \in [n]$*. Let* $\mathcal{J}$ *denote* $\gamma k \log \frac{k}{\delta}$ *samples with replacement from the discrete distribution* ρ *over* $[n]$*. If* $\left(V^T\right)_{:,\mathcal{J}} \Pi = QR$ *is a CPQR factorization, then there exists a permutation matrix* $\widehat{\Pi}$ *and a diagonal matrix* D *with* $D_{ii} = \pm 1$ *such that*

$$\|W\widehat{\Pi}D - VQ\|_F \leq \epsilon(1+\sqrt{2})k\sqrt{n} + \mathcal{O}\left(\epsilon^2\right).$$

Proof. The proof mirrors that of Theorem 1, we simply use Lemma 2 in place of Lemma 1 to ensure that the permutation $\widehat{\Pi}$ exists. □

Collectively, Theorems 1 and 2 tell us that if the subspace spanned by V_k is close to one that contains indicator vectors on the clusters, both of the proposed algorithms will yield an orthogonal transformation of V_k that approximates the indicators on the clusters. This representation of the subspace may then be directly used for cluster assignment. If desired, one may also use the computed clustering to seed the `k-means` algorithm and attempt to further improve the results with respect to the `k-means` objective function. This can either be accomplished by computing the cluster centers, or by using the points associated with the set $\mathcal{C}$ as the initial cluster centers.

3 Numerical results

We now present numerical simulations to validate the performance of our algorithm through examination of its behavior for multi-way spectral clustering. We consider the SBM with $k = 9$ equisized clusters and show via simulation that our algorithm, in contrast to `k-means++`, recovers the phase transition behavior at the information theoretic threshold [1, 2, 16, 3]. We also demonstrate similar phase transition behavior when there are $k = 7$ unequisized clusters.

For our tests we compare the behavior of three different clustering schemes: our algorithm, `k-means` using the columns $\mathcal{C}$ from our algorithm as an initial guess for cluster centers, and `k-means++` itself. All of our experiments were conducted using MATLAB® and the included `k-means++` implementation was used as a point of comparison. If `k-means` failed to converge after 100 iterations (the default maximum) we simply took the output as-is. We exclusively utilize the randomized variant of our algorithm with an oversampling factor of $\gamma = 4$. For more general problems, the non-random method may perform better, however, in this context they performed almost identically and hence the deterministic results are omitted.

Interestingly, in the regime where p and q are both $\Theta(\log n/n)$ we do not necessarily expect rapid enough asymptotic concentration of A for our results to theoretically guarantee recovery to the information theoretic limit. Nevertheless, we do observe good behavior of our algorithm when applied in this setting and, as we will see, the same behavior is not observed with `k-means++`.

3.1 Equisized clusters

First, we consider $k = 9$ equisized clusters each with size $m = 150$ (*i.e.*, $n = 1350$) and within-cluster and between-cluster connection probabilities $p = \alpha \log m/m$ and $q = \beta \log m/m$, respectively. Defining a grid of 80 equispaced values of α and 40 equispaced values of β, we generated 50 instances of the SBM for each (α, β) pair (redrawing the adjacency matrix if it was disconnected). Then, for both the adjacency matrix A and the degree-normalized adjacency matrix A_N with entries $(A_N)_{ij} = A_{ij}/\sqrt{d_i d_j}$, where d_i is the degree of node i, we computed the top nine eigenvectors and used them as input to each of the three algorithms mentioned before. While for degree-regular graphs we do not expect to observe differences in the results between A and A_N, when applied to graphs not coming from the SBM or that are not as degree-regular we anticipate better performance from degree-normalization.

In the plots of Figure 1 we color each point according to the fraction of trials resulting in successful recovery for each (α, β) pair, *i.e.*, the fraction of trials in which all nodes were correctly clustered. These phase diagrams show the primary advantage of our algorithm over `k-means++`: robust recovery (where possible). Whether used to explicitly compute the clustering directly, or to seed `k-means`, we cleanly recover a sharp phase transition that hews to the theory line. In contrast, `k-means++` fails to exactly recover the clusters a substantial portion of the time, even far from the phase transition.

Theoretically, if we scale k as $k = o(\log m)$ and take the asymptotic limit as $m \to \infty$ then a sharp phase transition between where recovery is possible almost surely and where it is not occurs at the threshold curve $\sqrt{\alpha} - \sqrt{\beta} = 1$ [3, 16]. In Figure 1 we observe that our phase transition does not quite match the asymptotically expected behavior. However, based on results for the two-block case [1], various arguments [3] show that we expect the location of the recovery threshold in the finite case to deviate from the asymptotic threshold with error that decays only very slowly, *i.e.*, $\Omega(1/\log m)$.

3.2 Unequisized clusters

Secondly, we consider the case where there are $k = 7$ *unequisized* clusters with sizes $m = 70, 80, 90, 100, 110, 120,$ and 130. In this case we define the within- and between-cluster connection probabilities as $p = \alpha \log 70/70$ and $q = \beta \log 70/70$. As before, we test the algorithms for exact recovery and present the phase plots in Figure 2. As before, our algorithm shows a sharp phase transition whereas `k-means++` does not, though we now lack any theoretical expectation for its exact location. Importantly, our algorithm seamlessly deals with unknown cluster sizes.

4 Discussion and conclusion

We have presented a new efficient (particularly the randomized variant) algorithm for spectral clustering of graphs with community structure. In contrast to the traditionally used `k-means` algorithm, our method requires no initial guess for cluster centers and achieves the theoretically expected recovery results for the SBM. Given that a bad initial guess can mean the `k-means` algorithm does not achieve the desired result this is a particularly important feature. Furthermore, we can always use our algorithm to generate an initial seed for `k-means` and observe in our experiments that this can provide small gains in recovery near the phase transition boundary.

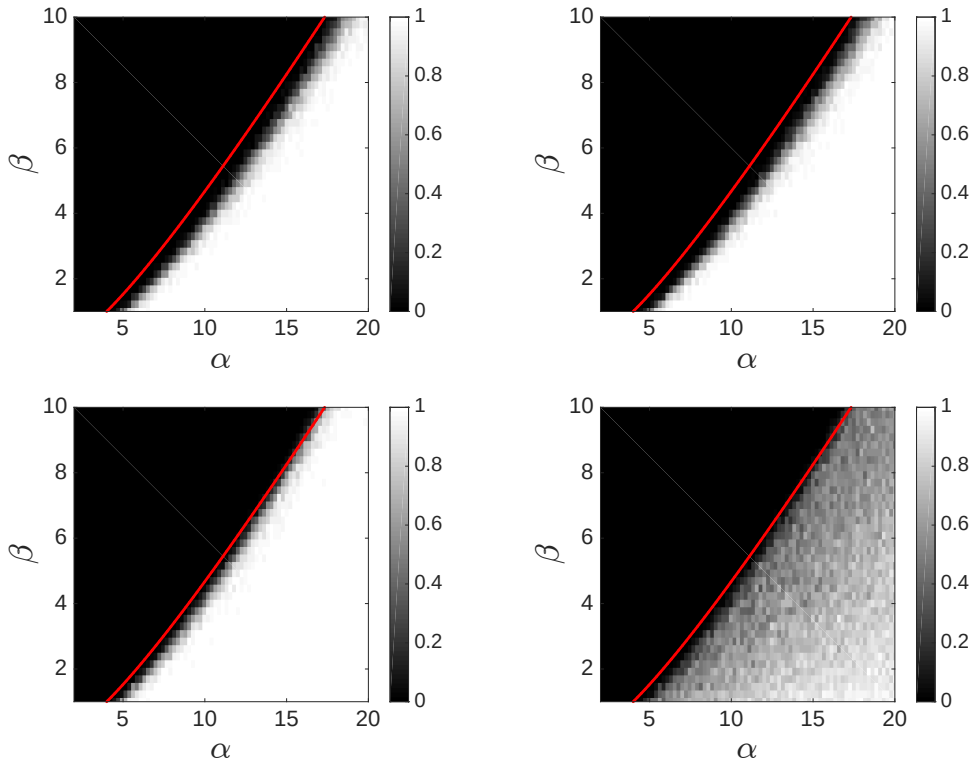

Figure 1: The phase plots here show the fraction of trials exhibiting exact recovery in the case of equisized clusters, and the asymptotic transition curve $\sqrt{\alpha} - \sqrt{\beta} = 1$ is given in red. The top row of phase plots corresponds to our randomized algorithm applied to eigenvectors V_k of A (top-left) and of A_N (top-right). We also give results based on seeding `k-means` with the clustering $\mathcal{C}$ from our algorithm applied to eigenvectors from A_N (bottom-left), and results for vanilla `k-means++` (bottom-right).

Recent results have yielded important understanding of what can be possible for recovery in the SBM. However, the SDP based methods that achieve these results do not scale to large problems of practical interest. Conversely, the traditionally used `k-means` algorithm, while scalable to large problems, fails to achieve the best possible behavior on these small scale SBMs due to its dependence on an initial guess. Our algorithm is both scalable to large problems and matches the behavior of SDPs on smaller model problems. These two properties make it attractive for general use.

Here we have explored the behavior of our algorithm on the SBM and provided theoretical justification for its use. To motivate its use for more general problems, we appeal to the expectation of similar orthogonal structure among the clusters. Furthermore, Theorems 1 and 2 may be extensible to cluster indicators with a more general structure, though at the expense of weaker results. We intend to explore the behavior of our algorithms on real-world graphs in future work.

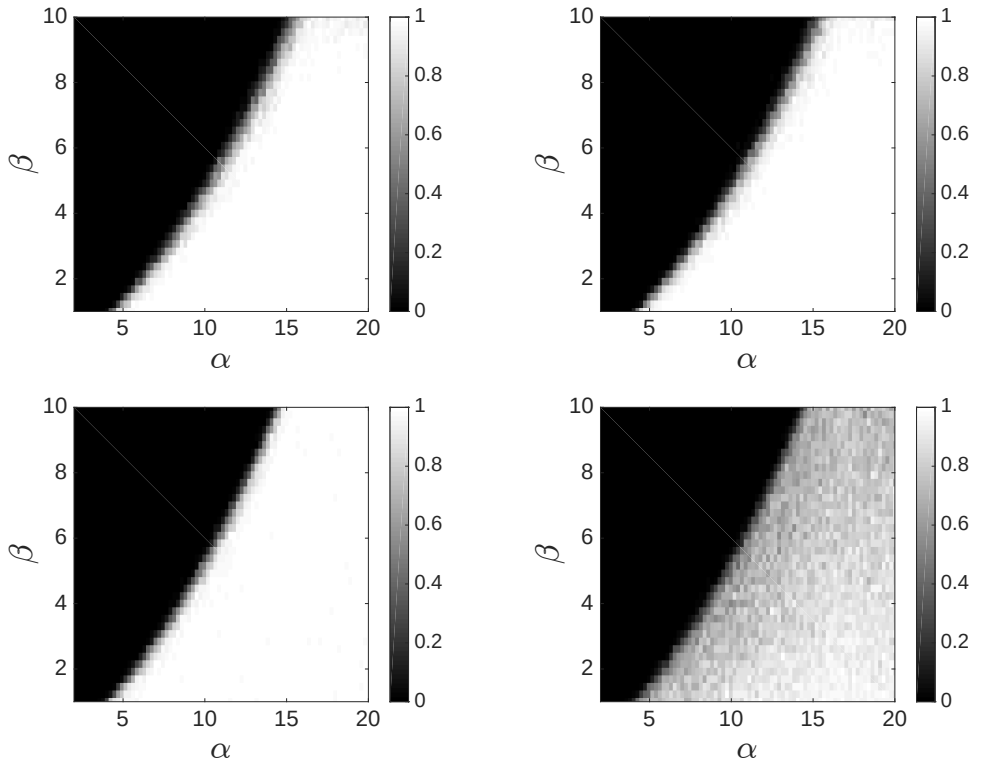

Figure 2: The phase plots here show the fraction of trials exhibiting exact recovery in the case of unequisized clusters. The top row of phase plots corresponds to our randomized algorithm applied to eigenvectors V_k of A (top-left) and of A_N (top-right). We also give results based on seeding `k-means` with the clustering $\mathcal{C}$ from our algorithm applied to eigenvectors from A_N (bottom-left), and results for vanilla `k-means++` (bottom-right).

5 Acknowledgments

A.D. is supported by a National Science Foundation Mathematical Sciences Postdoctoral Research Fellowship under grant number DMS-1606277. V.M. is supported by a Stanford Graduate Fellowship and a U.S. Department of Energy Computational Science Graduate Fellowship under grant number DE-FG02-97ER25308. L.Y. is supported by the U.S. Department of Energys Advanced Scientific Computing Research program under award DE-FC02-13ER26134/DE-SC0009409. The authors thank Austin Benson, Sven Schmit, and Nolan Skochdopole for useful discussion, as well as Stanford University and the Stanford Research Computing Center for providing computational resources and support that have contributed to these research results.

References

[1] E. Abbe, A. S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62(1):471–487, Jan 2016.

[2] Emmanuel Abbe and Colin Sandon. Community detection in general stochastic block models: Fundamental limits and efficient algorithms for recovery. In *Foundations of Computer Science (FOCS), 2015 IEEE 56th Annual Symposium on*, pages 670–688. IEEE, 2015.

[3] Naman Agarwal, Afonso S Bandeira, Konstantinos Koiliaris, and Alexandra Kolla. Multisection in the stochastic block model using semidefinite programming. *arXiv preprint arXiv:1507.02323*, 2015.

[4] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. SIAM, Philadelphia, PA, third edition, 1999.

[5] David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proc of the 18th annual ACM-SIAM symposium on discrete algorithms*, pages 1027–1035. SIAM, 2007.

[6] Peter Businger and Gene H Golub. Linear least squares solutions by householder transformations. *Numerische Mathematik*, 7(3):269–276, 1965.

[7] M Emre Celebi, Hassan A Kingravi, and Patricio A Vela. A comparative study of efficient initialization methods for the k-means clustering algorithm. *Expert Systems with Applications*, 40(1):200–210, 2013.

[8] Shivkumar Chandrasekaran and Ilse CF Ipsen. On rank-revealing factorisations. *SIAM Journal on Matrix Analysis and Applications*, 15(2):592–622, 1994.

[9] A. Damle, L. Lin, and L. Ying. Accelerating selected columns of the density matrix computations via approximate column selection. *ArXiv e-prints*, April 2016.

[10] Anil Damle, Lin Lin, and Lexing Ying. Compressed representation of KohnSham orbitals via selected columns of the density matrix. *Journal of Chemical Theory and Computation*, 11(4):1463–1469, 2015. PMID: 26574357.

[11] Chandler Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

[12] William E Donath and Alan J Hoffman. Lower bounds for the partitioning of graphs. *IBM Journal of Research and Development*, 17(5):420–425, 1973.

[13] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak mathematical journal*, 23(2):298–305, 1973.

[14] Gene H. Golub and Charles F. Van Loan. *Matrix Computations (3rd Ed.)*. Johns Hopkins University Press, Baltimore, MD, USA, 1996.

[15] Ming Gu and Stanley C Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM Journal on Scientific Computing*, 17(4):848–869, 1996.

[16] Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming: Extensions. *arXiv preprint arXiv:1502.07738*, 2015.

[17] Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Transactions on Information Theory*, 62(5):2788–2797, 2016.

[18] Paul W. Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109 – 137, 1983.

[19] Florent Krzakala, Cristopher Moore, Elchanan Mossel, Joe Neeman, Allan Sly, Lenka Zdeborová, and Pan Zhang. Spectral redemption in clustering sparse networks. *Proc Natl Acad Sci USA*, 110(52):20935–20940, 2013.

[20] S. Lloyd. Least squares quantization in pcm. *IEEE Trans. Inf. Theor.*, 28(2):129–137, September 2006.

[21] Michael W Mahoney and Petros Drineas. Cur matrix decompositions for improved data analysis. *Proc Natl Acad Sci USA*, 106(3):697–702, 2009.

[22] Laurent Massoulié. Community detection thresholds and the weak Ramanujan property. In *Proc of the 46th Annual ACM Symposium on Theory of Computing*, pages 694–703. ACM, 2014.

[23] Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. In *Proc of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2016, pages 814–827, New York, NY, USA, 2016. ACM.

[24] Rajeev Motwani and Prabhakar Raghavan. *Randomized algorithms*. Chapman & Hall/CRC, 2010.

[25] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Advances in Neural Information Processing Systems*, pages 849–856. MIT Press, 2001.

[26] Gregorio Quintana-Ort, Xiaobai Sun, and Christian H. Bischof. A blas-3 version of the qr factorization with column pivoting. *SIAM Journal on Scientific Computing*, 19(5):1486–1494, 1998.

[27] Karl Rohe, Sourav Chatterjee, and Bin Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *Ann. Statist.*, 39(4):1878–1915, 08 2011.

[28] G. W. Stewart. On the perturbation of LU, Cholesky, and QR factorizations. *SIAM Journal on Matrix Analysis and Applications*, 14(4):1141–1145, 1993.

[29] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015.

[30] Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007.

[31] Hongyuan Zha, Xiaofeng He, Chris Ding, Ming Gu, and Horst D Simon. Spectral relaxation for k-means clustering. In *Advances in neural information processing systems*, pages 1057–1064, 2001.